



4th Symposium of the Los Alamos Computer Science Institute: LACSI 2003—27 - 29 October

LACSI 2003 Chair: Dr. Rod Oldehoeft (rro@lanl.gov); **Program Chair:** Dr. Rob Fowler (rjf@rice.edu)

Symposium Site: Eldorado Hotel, 309 West San Francisco Street, Santa Fe, NM 87501

The Fourth Symposium offers participants a wide variety of opportunities to discuss and learn about technical and policy issues in high-performance computing.

Monday, October 27: Workshops on subjects of special interest; the Welcome Reception and Poster Exhibit are featured. Workshop room assignments will be posted at the Check-in/Registration desk.

REGISTRATION/CHECK-IN	South Concourse	8:00 am – 8:00 pm
CONTINENTAL BREAKFAST	South Concourse	8:00 – 9:00
Workshop (1) Clustermatic: A Revolutionary Approach (full day)	DeVargas	9:00 – 12:30
(2) System & Application Performance (full day)	Anasazi North	9:00 – 12:30
(3) Interoperability of Object Libraries (half day)	Zia A	9:00 – 12:30
(4) Mimetic Methods (half day)	Zia B	9:00 – 12:30
(5) High Availability & Performance Computing (half day)	Zia C	9:00 – 12:30
AM BREAK	South Concourse	10:30 - 11:00
LUNCH	(on your own)	12:30 – 2:00
Workshop (1) Clustermatic Computing (conclusion)	DeVargas	2:00 – 5:30
(2) System & Application Performance (conclusion)	Anasazi North	2:00 – 5:30
(6) Simulation-Driven Optimization (half day)	Zia A	2:00 – 5:30
(7) Particle Simulation (half day)	Zia B	2:00 – 5:30
(8) Adaptive and MPI Charm++ (half day)	Zia C	2:00 – 5:30
PM BREAK	South Concourse	3:30 – 4:00
Welcome Reception/Poster Exhibit	Anasazi Ballroom South	6:00 – 8:00

Tuesday, October 28: Important keynote address and refereed technical papers are presented by leading researchers in high-performance computing. All sessions will be held in the Anasazi Ballroom South. Registration/check-in will be in the South Concourse.

CONTINENTAL BREAKFAST	South Concourse	8:00 – 9:00
Welcome: Andy White, LANL, & Ken Kennedy, Rice University, LACSI Co-Directors		9:00 – 9:10
Keynote Address: <i>Challenges in Cyber Trust</i> , Dr. Eugene H. Spafford, Purdue University		9:10 – 10:00
AM BREAK	South Concourse	10:00 - 10:30
Reviewed Papers I: Partitioning, Distribution, and Scheduling		10:30 - 12:30
LUNCH	(on your own)	12:30 – 2:00
Reviewed Papers II: Compilation and Memory Hierarchies		2:00 – 3:30
PM BREAK	South Concourse	3:30 – 4:00
Reviewed Papers III: Cluster Communication and Storage		4:00 – 5:30

Wednesday, October 29: Four panel discussions focusing on important issues for researchers and managers in high-performance computing; the panel sessions will be held in Zia A, B, and C.

CONTINENTAL BREAKFAST	South Concourse	8:00 – 9:00
Panel I: The Earth Simulator: One Year Later		9:00 – 10:30
AM BREAK	South Concourse	10:30 - 11:00
Panel II: Experiences with TOP500 Linux Systems		11:00 – 12:30
LUNCH	(on your own)	12:30 – 2:00
Panel III: Phase II of DARPA's High-Productivity Computer Systems Program		2:00 – 3:30
PM BREAK	South Concourse	3:30 – 4:00
Panel IV: Federal Funding in High-Performance Computing		4:00 - 5:30
Closing Reception	Sunset Room	5:30 – 7:00

LACSI 2003: Keynote Address

(28 October 2003)

CHALLENGES IN CYBER TRUST

by

Dr. Eugene H. Spafford, Director

Center for Education and Research in Information Assurance and Security
Purdue University

The last 30+ years of computing have primarily been focused on getting computing to work, at least most of the time. Now we are employing computing equipment in everything from wristwatches to the power grid, from automobiles to air traffic control. It is a vital part of every critical infrastructure to society, as well as most of our research and entertainment industries. Thus, we now need to focus on making the computing reliable and trustworthy. The goal of computing faster may need to be tempered by the question (from another field): Unsafe at any speed? What are some of the ways in which large-scale computing can assist in this effort? Are the security challenges posed by grid and cluster computing unique? This talk will explore some ideas from this space, and encourage the audience to start thinking of their own examples.

***** NOTES *****

LACSI 2003: Workshops and Tutorials

FULL-DAY SESSIONS:

(1) Clustermatic: A Revolutionary Approach to Simplified Cluster Computing (full day)

Sung-Eun Choi, Ronald Minnich, Matthew Sottile, and Gregory Watson (Los Alamos National Laboratory)

Contact: Gregory Watson - gwatson@lanl.gov (505-665-0726)

Web: <http://www.clustermatic.org>

Abstract: Clustermatic is a revolutionary software architecture that redefines cluster computing at all levels from the BIOS to the parallel programming environment. The design maximizes performance and availability by achieving significant improvements in booting and application startup times, minimizing points of failure and vastly simplifying management and administration. It is suitable for use on a wide range of architectures, and has been deployed on tiny clusters containing only two diskless nodes, and on a 1024-node, nine-teraflop cluster at LANL. The tutorial combines detailed technical information about the design and operation of Clustermatic software with practical examples of deployment on a typical cluster system.

(2) System and Application Performance (full day)

Adolfy Hoisie (Los Alamos National Laboratory); Allen Malony (University of Oregon); John Mellor-Crummey (Rice University)

Contact: Adolfy Hoisie – hoisie@lanl.gov

Building extreme-scale parallel systems and applications that achieve scalable performance is difficult because system components (processors, memory hierarchies, networks) have become more difficult to use in pursuit of a high fraction of peak performance. Large systems with many components make failure inevitable, so fault tolerance is an integral part of performance. The workshop will address (1) aspects of performance and reliability, spanning system architecture, network, applications, and system software design; (2) methodologies for performance analysis and optimization including benchmarking, modeling, tools development, tuning and steering; and (3) performance of future systems on a scientific workload of interest to the audience.

HALF-DAY SESSIONS:

(3) Interoperability of Object Libraries for Scientific Computing (Monday, am)

William W. Symes and Anthony D. Padula (Rice University); Roscoe A. Bartlett (Sandia National Laboratories); Craig Rasmussen (Los Alamos National Laboratory)

Contact: William Symes – symes@rice.edu

(Abstract TBA)

(4) Mimetic Methods for Radiation Transport and Diffusion (Monday, am)

Jim Morel and Mikhail Shashkov (Los Alamos National Laboratory); Yuri Kuznetsov (University of Houston)

Contact: Mikhail Shashkov - shashkov@lanl.gov (505-667-4400)

Abstract: Radiation transport is a major component of the ASCI program. The memory and CPU time required for transport generally dominate all of the other physics in an ASCI calculation. The numerical requirements for transport discretization are extremely demanding because schemes must function well in totally different regimes. Because transport is so expensive, approximations based upon diffusion theory are often used in place of full transport theory. Although diffusion discretizations based upon standard finite elements are quite mature, the same cannot be said for mixed finite-element methods, which must generally be used in ASCI applications to ensure local conservation properties.

(5) High Availability and Performance Computing (Monday, am)

Stephen L. Scott (ORNL); Chokchai Leangsuksun (Louisiana Tech University)

Contact: Stephen L. Scott – scottsl@ornl.gov

Web: <http://www.cenit.latech.edu/hapcw2003/>

Abstract: High-availability (HA) computing has long played a critical role in industry mission-critical applications. On the other hand, high-performance computing (HPC) has equally been a significant enabler to the R&D community for their scientific discoveries. The combination of HA and HPC together will lead to even greater benefits to industrial, academic, and research communities. This workshop provides an important venue for discussing state-of-the-art and ongoing research and development in HAPC. In addition to the presentation of reviewed papers, the workshop will include a panel discussion on relevant topics.

(6) Simulation-Driven Optimization (Monday, pm)

Matthias Heinkenschloss, William W. Symes, Yin Zhang (Rice University)

Web: <http://www.caam.rice.edu/~heinken/lacsi/simandopt.html>

Abstract: The purpose of the workshop is to bring together researchers interested in various aspects of Simulation Driven Optimization and to discuss needs and challenges arising in the solution of such problems, as well as recent developments to meet them. Talks in this workshop describe (potential) optimization problems, outline challenges that need to be met for their solution, or present new optimization algorithm developments. This workshop complements the workshop organized by W. W. Symes (Rice U.) on "Interoperability of Object Libraries in Scientific Computing" (WS 3).

(7) Particle Simulation: Challenges and Prospects (Monday, pm)

Salman Habib (Los Alamos National Laboratory); Robert D. Ryne (Lawrence Berkeley National Laboratory)

Contact: Robert Ryne – RDRyne@lbl.gov

(Abstract TBA)

(8) Adaptive MPI and Charm++ (Monday, pm)

Laxmikant V. Kale and Orion Lawlor (UIUC)

Contact: kale@cs.uiuc.edu or olawlor@uiuc.edu

Web: http://charm.cs.uiuc.edu/LACSI_tutorial.shtml

Abstract: Charm++ and Adaptive MPI (AMPI) have recently emerged as powerful parallel programming systems. NAMD, a Charm++ application, was a Gordon Bell Award-winning application in 2002. It achieved unprecedented speedup for molecular dynamics on 3000 processors of the PSC Lemieux. Other applications, such as *ab initio* MD, Cosmology and Rocket Simulation, have demonstrated that Charm++ is a viable and successful parallel programming model. This success stems from Charm++'s adaptive run-time system and processor virtualization. The idea of virtualization is simple: let the programmer divide the work into a large number of chunks independent of the number of processors and let the run-time system map these chunks to processors. This leads to an effective separation of concerns between the programmer and the runtime system. AMPI is an implementation of the MPI-1.1 Standard based on the Charm++ RTS. With virtualization of MPI processors, AMPI supports such features as automatic adaptive overlap of communication and computation, automatic checkpointing with restart, and automatic load balancing.

LACSI 2003: Poster Exhibits

- (1) **Connection Management for Massively Scalable TCP**
Patricia Eileen Crowley, Arthur B. Maccabe, and Breanne A. Duncan (The University of New Mexico)
- (2) **Grid Performance Prediction with Performance Skeletons**
Sukhdeep Sodhi and Jaspal Subhlok (University of Houston)
- (3) **A Portable High Performance Co-Array Fortran Compiler**
Yuri Dotsenko, Cristian Coarfa, and John Mellor-Crummey (Rice University)
- (4) **An In-House Development for Visualization in Molecular Science**
Jean M. Favre and Mario Valle (Swiss Center for Scientific Computing)
- (5) **Dragon: Program Analysis and Development Support for Open MP Codes**
Barbara Chapman, Oscar Hernandez, Lei Huang, Yi Wen, Tien-hsiung Weng, and Zhenying Liu (University of Houston)
- (6) **A Sample-Driven Call Stack Profiler**
Nathan Froyd, John Mellor-Crummey, and Nathan Tallent (Rice University)
- (7) **Clam: Light Communication Layer for Asynchronous Mobile Computations**
Andriy Fedorov and Nikos Chrisochoides (College of William and Mary)
- (8) **PAPI Version 3**
Jack Dongarra, Kevin London, Shirley Moore, Philip Mucci, Daniel Terpstra, Haihang You, and Min Zhou (University of Tennessee)
- (9) **Processor-Level Partitioning of Kernel and User Mode Functionality on SMP Systems**
Carl A. Sylvia (The University of New Mexico)
- (10) **Model Coupling Toolkit Benchmarks**
J. Walter Larson, Robert L. Jacob, and Everest T. Ong (Argonne National Laboratory)
- (11) **Fault Injection into MPI Programs**
Charng-da Lu, Karthik Pattabiraman, and Daniel A. Reed (University of Illinois at Urbana-Champaign)
- (12) **High Performance Object-Oriented Java Programming: Fiction or Reality?**
Zoran Budimlic (Rice University); Giovanni Lapenta (LANL), Stefano Markidis (Politecnico di Torino); and Brian VenderHeyden (Los Alamos National Laboratory)
- (13) **Profile-Based Dynamic Voltage Scaling for I/O-Intensive Codes**
Karthik Pattabiraman (Los Alamos National Laboratory and University of Illinois at Urbana-Champaign); Wu-Chun Feng (Los Alamos National Laboratory); and Daniel Reed (University of Illinois at Urbana-Champaign)
- (14) **A High-Level Approach to the Synthesis of High-Performance Codes for Quantum Chemistry**
David E Bernholdt, Venkatesh Choppella, and Robert J. Harrison (Oak Ridge National Laboratory); Alexander Auer and Marcel Nooijen (University of Waterloo); Gerald Baumgartner, Alina Bibireata, Daniel Cociorva, Xioyang Gao, Sriram Krishnamoorthy, Sandhya Krishnan, Chi-Chung Lam, Qingda Lu, Russell M. Pitzer, P. Sadayappan, and Alexander Sibiryakov (Ohio State University); So Hirata (Pacific Northwest National Laboratory); J. Ramanujam (Louisiana State University)
- (15) **Optimizing Collective Multicast on PSC Lemieux**
Sameer Kumar and Laxmikant V. Kale (University of Illinois at Urbana-Champaign)
- (16) **Supporting Generic Programming in a Multi-Language Component-Based Environment**
Wael R. Elwasif, Thomas C. Schulthess, and David E. Bernholdt (Oak Ridge National Laboratory); Gregory Brown (Florida State University)
- (17) **Design and Implementation of the Prophecy Automated Model Builder for Parallel and Grid Applications**
Joseph Paris (Northwestern University), Xingfu Wu and Valerie Taylor (Texas A&M University)

LACSI 2003: Technical Paper Sessions

Reviewed Papers I: Partitioning, Distribution, and Scheduling

- *Parallel Remote Method Invocation and M-by-N Data Redistribution*
Kostadin Damevski and Steven Parker (University of Utah)
- *An Approach to Parallel MxN Communication*
Felipe Bertrand, Yongquan Yuan, Kenneth Chiu, and Randall Bramley (Indiana University)
- *A Partitioner-Centric Model for SAMR Partitioning Trade-off Optimization: Part I*
Johan Steensland and Jaideep Ray (Sandia National Laboratories)
- *Opportune Job Shredding: An Effective Approach for Scheduling Parameter Sweep Applications*
Rohan Kurian, Pavan Balaji, and P. Sadayappan (The Ohio State University)

Reviewed Papers II: Compilation and Memory Hierarchies

- *Investigating Adaptive Compilation Using the MIPSpro Compiler*
Keith D. Cooper and Todd Waterman (Rice University)
- *On Reducing Storage Requirement of Scientific Applications*
Guohua Jin, and John Mellor-Crummey (Rice University)
- *Regression-Based Multi-Model Prediction of Data Reuse Signature*
Xipeng Shen, Yutao Zhong, and Chen Ding (University of Rochester)

Reviewed Papers III: Cluster Communication and Storage

- *Implementing Scalable Diskless Clusters Using NFS*
James Howard Laros III and Lee Ward (Sandia National Laboratories)
- *The LAM/MPI Checkpoint/Restart Framework: System-Initiated Checkpointing*
Sriram Sankaran, Jeffrey M. Squyres, Brian Barrett, Andrew Lumsdaine, Jason Duell, Paul Hargrove, and Eric Roman (Indiana University and Lawrence Berkeley National Laboratory)
- *Fault Tolerant Communication Library and Applications for High Performance Computing*
Graham Edward Fagg, Edgar Gabriel, Zizhon Chen, Thara Angskun, George Bosilica, Antonin Bukovsky, and Jack Dongarra (ICL/University of Tennessee, Knoxville)
- *High Performance Broadcast Support in LA-MPI over Quadrics*
Weikuan Yu, Sayantan Sur, and D.K. Panda (The Ohio State University); Rob T. Aulwes and Rich L. Graham (Los Alamos National Laboratory)

LACSI 2003: Panel Sessions

Panel I: The Earth Simulator: One Year Later

Discussants will review aspects of hardware, software, and applications of the system that has caused much controversy since it was first introduced in 2002.

Panel II: Experiences with TOP500 Linux Systems

Systems running Linux have made steady progress in invading the high end of the TOP500 list. Representatives from organizations with the most powerful systems will discuss their experiences.

Panel III: High Productivity Computer Systems, Phase II

Representatives from Cray, IBM, and Sun will discuss their efforts in Phase II of DARPA's HPCS program.

Panel IV: Federal Programs in High-Performance Computing

Panelists from Federal agencies that support HPC and that have participated in HECRTF will present perspectives on the future of Federal programs that support the research, development, and use of high performance computing platforms.